Biased gradient squared descent saddle point finding method

Juliana Duncan,[1] Qiliang Wu,[2] Keith Promislow,[2] and Graeme Henkelman[1]
[1]*Department of Chemistry and the Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, Texas 78712-0165, USA*
[2]*Department of Mathematics, Michigan State University, East Lansing, Michigan 48824, USA*

(Received 26 March 2014; accepted 28 April 2014; published online 15 May 2014)

The harmonic approximation to transition state theory simplifies the problem of calculating a chemical reaction rate to identifying relevant low energy saddle points in a chemical system. Here, we present a saddle point finding method which does not require knowledge of specific product states. In the method, the potential energy landscape is transformed into the square of the gradient, which converts all critical points of the original potential energy surface into global minima. A biasing term is added to the gradient squared landscape to stabilize the low energy saddle points near a minimum of interest, and destabilize other critical points. We demonstrate that this method is competitive with the dimer min-mode following method in terms of the number of force evaluations required to find a set of low-energy saddle points around a reactant minimum. © *2014 AIP Publishing LLC*. [http://dx.doi.org/10.1063/1.4875477]

I. INTRODUCTION

One of the major challenges in computational chemistry and condensed matter physics is the efficient calculation of the rate of chemical reactions and diffusion events which have barriers much higher than the thermal energy. Molecular dynamics (MD) is an important technique that can be used to calculate rates of reactions, but it is generally limited to simulation time scales of nanoseconds. Alternative methodology is needed to capture the long time scale dynamical evolution.

A powerful framework for calculating rates of chemical systems is transition state theory (TST). Within TST, the rate is estimated by the equilibrium flux through a hyper-surface separating a reactant from product states. The TST rate can be expressed as

$$k^{\mathrm{TST}} = \frac{1}{2}\langle |v|\delta(x - x^{\ddagger})\rangle_R, \tag{1}$$

where $x = x^{\ddagger}$ is the TST dividing surface, v is the velocity through the dividing surface, and $\langle \dots \rangle_R$ is the canonical average over the reactant state.[1,2] Determining the TST dividing surface can be the most difficult aspect of calculating a TST rate in high dimensional systems. Harmonic transition state theory (HTST) is a widely used simplification in which the dividing surface is approximated as a set of hyperplanes passing through each saddle point (SP) connecting the minimum to product states along minimum energy paths.[3,4] The normal to these hyperplanes are along the negative curvature modes of the potential at each SP. Expanding the potential up to second order around the minimum and saddles allows for an analytic evaluation of Eq. (1) to give the HTST rate expression,

$$k^{\mathrm{HTST}} = \frac{\prod_{i=1}^{3N} \nu_i^{\mathrm{init}}}{\prod_{i=1}^{3N-1} \nu_i^{\ddagger}} e^{-(E^{\ddagger} - E^{\mathrm{init}})/k_b T}, \tag{2}$$

where ν^{init} and $\nu^{\ddagger}$ are the normal mode frequencies and E^{init} and $E^{\ddagger}$ are the potential energies at the minimum and SP, respectively.[4] With such a simple rate expression, the hardest part of evaluating HTST rates is the determination of the SPs connecting the minimum to products via minimum energy pathways.

There are two classes of SP finding methods. The first class requires the product states as input for finding minimum energy paths and SPs on these paths. Such methods include the nudged-elastic band method,[5–7] and the closely related string method.[8] The second class uses only information about the location of the starting minimum. Popular methods include the min-mode following methods, the dimer method,[9–11] Rayleigh-Ritz minimization[12] as in the hybrid eigenvector following method,[13] or the Lanczos algorithm as in the activation relaxation technique (ART-nouveau),[14] to determine the lowest mode of the Hessian. A review of saddle point finding methods of both classes can be found in Refs. 15–17.

Methods that walk from a reactant minimum to a SP are important since these methods do not require advance knowledge of product states. An ideal SP finding method for estimating rates with HTST would only use information from the minimum and efficiently find a set of low-lying saddle points in the neighborhood of the reactant state. Partially motivated by the functionalized energies,[18,19] we introduce a SP finding method, the biased gradient squared descent (BGSD), that does not require calculating a minimal eigenvector and can be parametrically tuned to find low energy SPs.

II. METHODS

A. Biased gradient squared descent

Given a potential energy surface, $V(r)$, the goal of the BGSD method is to produce a family of associated energy landscapes, $H(r; \alpha, \beta)$ whose critical points include the critical points (CPs) of V, but with the property that selective tuning of the parameters α and β allow a systematic

transformation of the SPs of V into local minima of H. Such an algorithm would simplify the search for SPs of V into a search for local minima of H, a task ideally suited to gradient flow. We desire a construction of H which is readily generalizable, easy to implement and computationally efficient, and which produces as few extraneous local minima (local minima of H which are not also CPs of V) as possible. Moreover it is desirable to be able to selectively search for the saddle points of V, starting from those with an energy of a local minimum and moving upwards.

The starting point for the construction of the BGSD potential, H, is the squared gradient of V, $|\nabla V|^2$. Every CP of V is a degenerate global minimum of $|\nabla V|^2$ and all CPs of $|\nabla V|^2$ will either be CPs or inflection points (IPs) of V. It is natural to ask if this degeneracy can be systematically broken, mapping CPs of V with desired properties to local minima of H while the remainder of the V-CPs become H-SPs. One approach is to bias $|\nabla V|^2$ towards a level-set of V

$$H(r;\alpha,\beta) := \frac{1}{2}|\nabla V(r)|^2 + \frac{1}{2}\alpha(V(r)-\beta)^2, \qquad (3)$$

where β is the chosen level-set and α determines the bias strength. The CPs of H are the solutions of

$$\nabla H(r) = [\nabla^2 V(r) + \alpha(V(r)-\beta)]\nabla V(r) = 0, \qquad (4)$$

which fall into two classes, the CPs of V, that are solutions of $\nabla V(r) = 0$, and the "eigen-critical points" (EPs) r_{EP} for which $\nabla V(r_{\mathrm{EP}})$ is a non-zero eigenmode of the Hessian matrix $\nabla^2 V(r_{\mathrm{EP}})$ associated to eigenvalue $\lambda = -\alpha(V(r_{\mathrm{EP}})-\beta)$. Figure 1 illustrates a mapping from V to H for a one-dimensional potential, and highlights how CPs and IPs of V are mapped onto $|\nabla V|^2$ as extrema. The harmonic bias term in H stabilizes the extrema of $|\nabla V|^2$ whose energies are close to

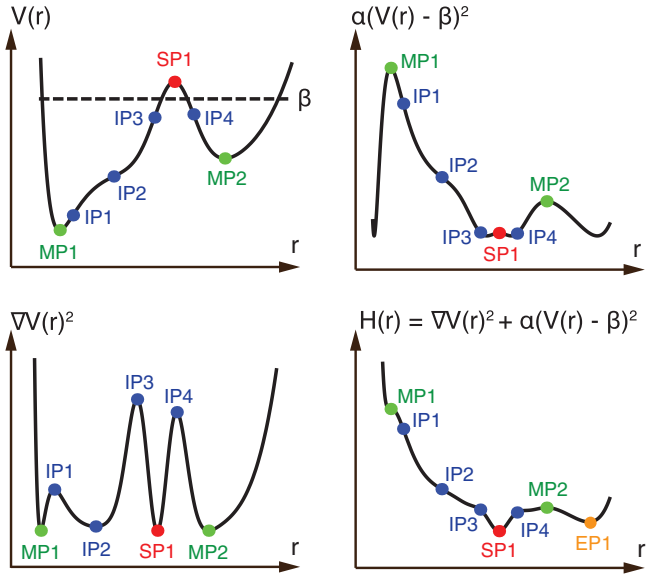

FIG. 1. Illustration of the BGSD method, depicting the transformation of the V- into the H-landscape. The plot shows the CPs of V on the different landscapes include the two local minima, MP1–2, the local maximum (representing a SP in a higher dimensional space), SP1, and the inflection points, IP1–4. This example shows how inflection points become local extrema in $|\nabla V|^2$. One eigen-critical point, EP1, is created in H due to the destabilization of MP2.

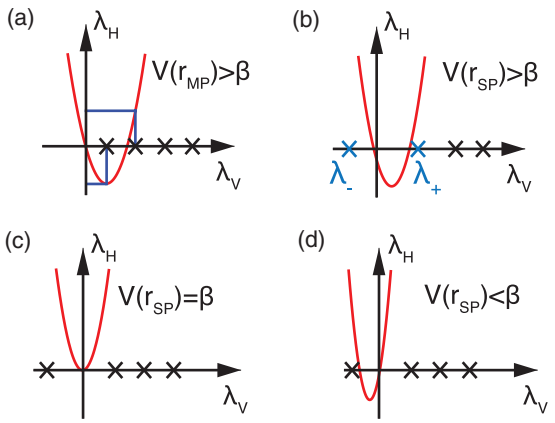

FIG. 2. Spectral mapping of the eigenvalues of the V- to the H-Hessian under the following conditions: (a) at the minimum, where one eigenvalue of the V-Hessian becomes negative in the H-Hessian; (b) at SPs with values of β that are greater than; (c) equal to; and (d) and less than $V(r_{\mathrm{SP}})$, but inside the H-stability window. Case (b) highlights the location of $\nabla^2 V$ eigenvalues, λ_- and λ_+, which define the H-stability window.

β. The BGSD method is similar to the gradient norm method for finding transition states[20] with the addition of a bias potential, commonly used in umbrella sampling methods,[21] which stabilizes critical points around a particular energy.

B. Selection of *H*

In a SP search about a local minimum (MP), r_{MP} of V, it is desirable to have r_{MP} map to a SP of H, and hence be unstable under an H-gradient flow, while the V-SPs should map to local minima of H. The stability of V-CPs, r_{CP}, under the H-flow, is determined by the H-Hessian, which takes a simple form at V-CPs,

$$\nabla^2 H(r_{\mathrm{CP}}) = [\nabla^2 V(r_{\mathrm{CP}}) + \alpha(V(r_{\mathrm{CP}})-\beta)]\nabla^2 V(r_{\mathrm{CP}}). \qquad (5)$$

In particular, for fixed values of the scalar $\alpha(V(r_{\mathrm{CP}})-\beta)$, Eq. (5) is a quadratic-matrix map on $\nabla^2 V$, so that each eigenvalue λ_V of $\nabla^2 V(r_{\mathrm{CP}})$ is mapped onto an associated eigenvalue λ_H of $\nabla^2 H$ via the relation

$$\lambda_H = \lambda_V^2 + \alpha(V(r_{\mathrm{CP}})-\beta)\lambda_V =: g(\lambda_V; V(r_{\mathrm{CP}})). \qquad (6)$$

As depicted in Fig. 2, this relation is the basis for the spectral mapping of CPs of V onto saddles or local minima of H. The sign of $(V(r_{\mathrm{CP}})-\beta)$ plays an essential role in determining the stability of V-CPs under the H-gradient flow. For any CP, r_{CP} of V, the associated parabola g always passes through the origin. However, if $V(r_{\mathrm{CP}})-\beta < 0$ then g is negative on an interval to the right of the origin. In particular at r_{MP}, the spectrum of $\nabla^2 V(r_{\mathrm{MP}})$ is positive and choosing α sufficiently large, the V-MP maps to an H-SP. However, if r_{CP} is a first order SP of V, then $\nabla^2 V(r_{\mathrm{SP}})$ has a single negative eigenvalue, and sweeping the value of β from below $V(r_{\mathrm{SP}})$ to above, the negative region of the parabola g, initially on the positive real axis, shrinks as β approaches $V(r_{\mathrm{SP}})$, compresses to a point when $\beta = V(r_{\mathrm{SP}})$ and then switches to the negative real axis for $\beta > V(r_{\mathrm{SP}})$, and finally grows to encompass the negative eigenvalue of $\nabla^2 V(r_{\mathrm{SP}})$. The key to tuning H is that r_{CP} will be an H-SP when $|V(r_{\mathrm{CP}})-\beta|$ is

sufficiently large, and an H-MP when $|V(r_{CP}) - \beta|$ is sufficiently small, so long as the Hessian $\nabla^2 V(r_{CP})$ has no zero eigenvalues. Indeed, r_{CP} is an H-MP for beta satisfying $-\lambda_- < \alpha(V(r_{CP}) - \beta) < -\lambda_+$, where λ_- is the negative eigenvalue of $\nabla^2 V(r_{CP})$ with the smallest absolute magnitude and λ_+ is the positive eigenvalue with the smallest absolute magnitude. We call these values of β the H-stability window.

The optimal choice for α depends upon the stiffness of the V energy landscape. If $\nabla^2 V$ has many small eigenvalues, then α should be correspondingly small so that reasonable increments in β will not jump over the H-stability window. However if the smallest eigenvalues of $\nabla^2 V$ are large, then α should be chosen appropriately large so that the H-stability intervals of differing V-CPs do not overlap excessively.

C. Stability and bifurcation of eigen-critical points

The H-stability of the EPs does not admit a general analysis, however one can understand their stability by the local bifurcations that create them. Indeed, for $\alpha = 0$ there are no EPs of H, and as α is increased from zero the EPs will typically be created in bifurcations from V-CPs (see Fig. 3). Returning to the example of the first order V-CP, r_{CP}, depicted in Fig. 2, for a fixed value of $\alpha > 0$, as β is varied, an eigenvalue $\lambda_H(r_{CP}; \alpha, \beta)$ of $\nabla^2 H(r_{CP})$ will pass from positive to negative as β leaves the H-stability interval for r_{CP}. The value of β for which $\lambda_H(r_{CP}; \alpha, \beta) = 0$ corresponds to the creation of a EP, r_{EP} of H. In particular r_{EP} will exist when $\lambda_H(r_{CP}, \alpha, \beta) < 0$ and will move away from r_{CP} as $\lambda_H(r_{CP}; \alpha, \beta)$ decreases from zero. For β sufficiently close to the critical value at which r_{EP} is created, then r_{EP} will be a local minimum of H, but it may become unstable to secondary bifurcations as β is further increased. If an EP is sufficiently close to r_{CP}, then a gradient flow may converge to r_{EP}, however we can subsequently adjust α smaller, perhaps to zero, which will eliminate the EP and re-stabilize r_{CP}, so that further evolution of the H-gradient flow for the $\alpha = 0$ may converge to r_{CP}. It is important to note that minimization of the $H(r; 0, \beta) = |\nabla V|^2$ landscape does not guarantee convergence to a CP of V, since IPs of V can be local minima (see Fig. 1).

D. Implementation of *H* gradient flow

The SP search is achieved by descending to minima of H through an H gradient flow,

$$\frac{dr}{dt} = -\nabla H(r), \tag{7}$$

where ∇H is given in Eq. (4). An efficient implementation of this flow must avoid a full calculation of the Hessian of V, which is generically excessively large. However we require only the quantity $\nabla^2 V(r)\nabla V(r)$, which admits an efficient matrix-free evaluation through the difference quotient

$$\nabla^2 V(r)\nabla V(r) = \lim_{\delta \to 0} \frac{\nabla V\,[r + \delta \nabla V(r)] - \nabla V(r)}{\delta}. \tag{8}$$

For an appropriately small choice of δ, this different quotient yields an accurate approximation of $\nabla^2 V(r)\nabla V(r)$ at the cost of two force evaluations.

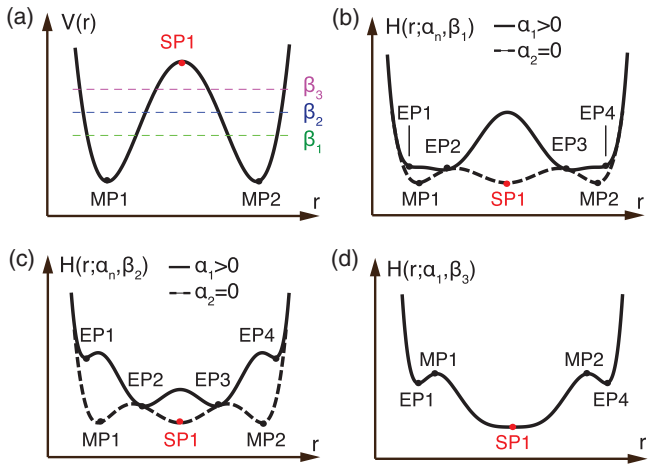

FIG. 3. Illustration of the bifurcations of EPs in BGSD. (a) $V(r)$ with the CPs highlighted and the three values of β used in (b), (c), and (d). Two values of α are shown above, a fixed $\alpha_1 > 0$ and $\alpha_2 = 0$ which eliminates the EPs in H. (b) In the $H(r; \alpha_n, \beta_1)$ landscape, where $\lambda(r_{CP}; \alpha_1, \beta_1) < 0$, two pairs of EPs are formed, {EP1, EP4}, and {EP2, EP3}, the H-gradient flow will converge to an EP, and subsequent evolution to $\alpha = \alpha_2$ will yield the return to the V-MP. (c) In the $H(r; \alpha_n, \beta_2,)$ landscape, we still have $\lambda(r_{CP}; \alpha_1, \beta_2) < 0$, however, H-gradient flow which converges to EP2 and EP3 will subsequently converge to SP1 with $\alpha = \alpha_2$. (d) In the $H(r; \alpha, \beta_3)$ landscape, $\lambda(r_{CP}; \alpha, \beta_3) > 0$, so SP1 is stabilized and absorbs EP2 and EP3. Notice that EP1 and EP4 are formed along directions in the potential which do not lead to SPs.

As discussed previously, adjusting α to smaller values will eliminate EPs of H while minimization of $|\nabla V|^2$ eliminates all EPs of H. The parameter α is adjusted by first using a fixed $\alpha > 0$. After convergence of $H(r; \alpha, \beta)$, the resulting r is used as the initial data point in an optimization of the $H(r; 0, \beta)$ landscape. The value of $H(r; 0, \beta)$ can be used as a simple check for convergence to IPs. All CPs of V are zero on the $|\nabla V|^2$ landscape, so convergence to a positive value of $|\nabla V|^2$ indicates that an IP has been found.

III. RESULTS

A. LEPS potential in two dimensions

As a simple two dimensional test of the BGSD method, a London-Eyring-Polanyi-Sato (LEPS) potential coupled to a harmonic oscillator (described elsewhere[22]) was chosen to illustrate the issues that arise in BGSD SP searches. Two Gaussian functions were added to this potential to increase the number of SPs.[23] The potential is shown in Fig. 4(a); it has four local minima, MP1–4, four first order SPs, SP1–4, and one maximum, SSP1. The maximum has two negative modes and is thus labelled as a second order saddle point (SSP), to make a connection with the higher dimensional systems described subsequently. MP1 is taken as the reactant state minimum. SP1, SP3, and SP4 are connected to MP1 by a steepest descent path; SP2 is not.

In the gradient squared landscape, shown in Fig. 4(b), all CPs of V are global minima with a value of zero. One important feature of this landscape is the low curvature direction near MP1, which does not lead to a CP of V. The biasing term, $\frac{1}{2}\alpha(V(r) - \beta)^2$, with $\alpha = 5$ is added to the gradient squared landscape to obtain $H(r; \alpha, \beta_1)$, as shown in Fig. 4(c). A value

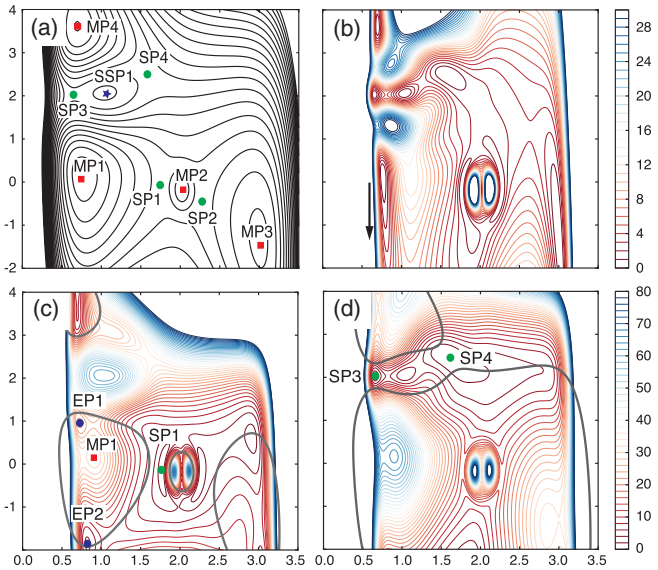

FIG. 4. (a) Contour of the modified LEPS potential, $V(r)$, with the CPs labelled. (b) Contour plot of $|\nabla V(r)|^2$. The black arrow is parallel to a soft mode which does not lead to a CP of $V(r)$. (c) Contour plot of H with a value of $\beta_1 = 2.5$, less than SP1, and (d) $\beta_2 = 5.0$, less than SP4. The grey lines indicate the isosurface where $V(r) = \beta$.

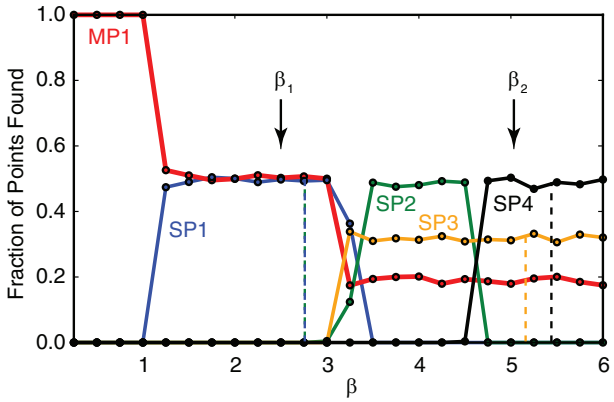

FIG. 5. CPs of the LEPS potential found as a function of β. The vertical dashed lines indicate the energy of the CP of the corresponding color. The marked values of β correspond to Figs. 4(c) and 4(d).

of $\beta_1 = 2.5$ was chosen to be slightly less than the lowest energy saddle, SP1. To follow gradient descent trajectories of H, initial points were drawn from a Monte Carlo (MC) sampling of the contour $V(r) = \beta$, shown as grey lines in Fig. 4(c). The MC sampling was initiated from MP1 and hence only sampled points from the branch of the $V(r) = \beta$ around MP1.

Three points were found by minimizing $H(r; \alpha, \beta_1)$ from initial points on β_1: EP1, EP2, and SP1. There are other minima on the H landscape, but the initial points, drawn from the branch of the β_1 contour around MP1, do not lie in their basin of attraction. The three minima of H are from different classes. SP1 is a saddle of V, and a target of the search. EP1 is along the path to the higher energy saddle, SP3. EP2 is along the soft mode for which there is no CP of V. Subsequent minimization of $H(r; 0, \beta_1) = |\nabla V|^2$ takes points at both EP1 and EP2 back to the minimum, MP1. At the higher value of β_2, shown in Fig. 4(d), minimization of $H(r; \alpha, \beta_2)$ finds the higher energy saddles SP3 and SP4.

A more systematic search of saddles around the reactant is achieved by scanning over values of β, increasing from the energy of the minimum. In Fig. 5, 24 values of β were chosen between the energy of the minimum MP1 and the maximum SSP1. A total of 500 searches were used to determine the fraction of CPs found at each value of β. For $\beta > 1.0$ the H-stability window for SP1 opens and there is a sharp increase in the number of times this CP is found. Above the energy of SP1, the β contour extends beyond the reactant state basin so that for $\beta > 3.0$, SP2, which is not connected with MP1, is frequently found. SP3 reaches its H-stability window at this energy as well and is found over a wide range of β. SP4 is stable for $\beta > 4.5$ and is found frequently due to a large basin of attraction as compared to SP3. MP1 is found for all values of β due to the minimum in $H(r; \alpha, \beta)$ along the soft mode (e.g., EP2), which relaxes back to MP1 upon minimization of $H(r; 0, \beta)$. At the highest values of β, it is possible to find the maximum, SSP1, but it is sandwiched between SP3 and SP4 and has a small basin of attraction.

These results demonstrate that the BGSD method can find all SPs for our two-dimensional system. The frequency and the values of β that CPs are found depend on the size and location of the basin of attraction. Two weaknesses of the method are identified. The first is that when β exceeds the energy of the lowest saddle, the BGSD can find saddles that are not connected to the initial state minimum. Second, soft modes introduce EPs in H which relax back to the minimum under $|\nabla V(r)|^2$. This later issue is worrisome for higher dimensional systems which typically have many soft modes.

B. Al adatom on an Al(100) surface

A more challenging problem is the search for diffusion mechanisms of an Al adatom on the surface of Al(100). The Al interatomic potential was taken to be of the embedded atom method form from Voter and Chen.[24] The surface was modeled as six layers containing 64 atoms with a single adatom on the surface. The bottom two layers were held frozen in bulk positions, leaving 771 degrees of freedom. The six saddle points within 0.5 eV of the reactant minimum (excluding those equivalent by symmetry) are shown in Fig. 6. The two lowest energy diffusion mechanisms are the exchange (SP1) and hop (SP2). SP3 involves four surface atoms. SP4 and SP5 are longer range exchange processes involving three and two surface atoms, respectively. Additionally, there is a second order saddle, SSP1, which was reported by Maronsson *et al.*[25]

A concern with the implementation of the BGSD method is that EPs created by soft modes will dominate the minima on the H landscape as the dimensionality of the system increases. A dimensionality scaling test was performed by changing the number of atoms fixed to their equilibrium positions. The lowest dimensional case had the bottom five atomic layers fixed, leaving 195 free degrees of freedom. The highest dimensional system allowed all atoms to move, corresponding to 1155

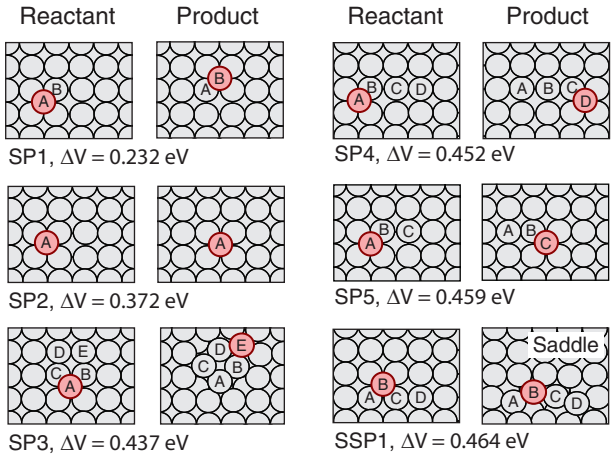

FIG. 6. Five diffusion mechanisms with barriers less than 0.5 eV for the adatom on an Al(100) surface, as well as one second order saddle point.

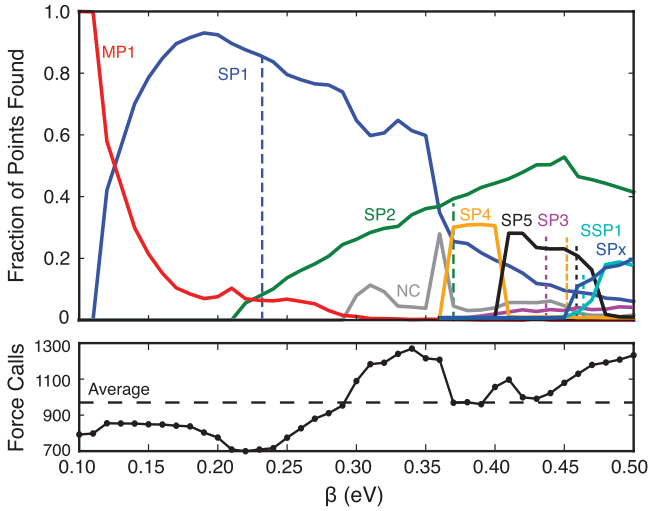

FIG. 8. Distribution of CPs found in the Al adatom system as a function of the bias energy, β. The low energy CPs are shown in Fig. 6. SPx indicates any higher energy saddles found, and NC are any CPs that are not connected to the reactant, MP1.

degrees of freedom. To test for the number of EP found involving soft modes, 200 minimizations of H were performed on each system. A fixed value of $\alpha = 10$ was used, and β was set to the lowest SP energy for each system, which varied somewhat with the number of frozen atoms. From each initial point on the β contour, $H(r; \alpha, \beta)$ was first minimized followed by $H(r; 0, \beta) = |\nabla V(r)|^2$. The CPs found were either the reactant minimum, or a SP connected to the reactant state. As the dimensionality was increased, the fraction of SPs decreased and the number of EPs along soft modes increased, as shown in Fig. 7 (blue line).

To help solve this problem, local displacements were made from the minimum to the β contour. Specifically, only the adatom and all atoms within 3.3 Å were displaced. The number of SPs found was significantly improved using this local displacement scheme, as shown in Fig. 7 (red line), as well as the scaling with dimensionality. The local biasing scheme was used in all other results for the Al system.

The fraction of different saddles found was calculated over a set of β values, as shown in Fig. 8. At low values of β, only the minimum (MP1) is found and then as β increases, this fraction drops to zero. Convergence to different saddle points, with increasing β, occurs roughly in the order of their energies. There are some exceptions, but Fig. 8 clearly illustrates the turning on and off as β passes through the respective H-stability windows of the different saddles points. At high values of β, saddles with energies larger than 0.5 eV are found (SPx), as well as the second order saddle (including SSP1), SPs that are not connected (NC) to MP1, and very few IPs.

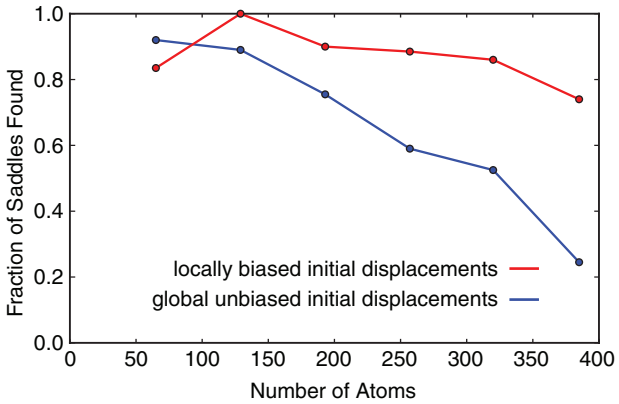

FIG. 7. The BGSD method scales poorly with dimensionality when searches are initiated by displacing all degrees of freedom from a minimum in the Al adatom system; local displacements improve the scaling by avoiding convergence to spurious EPs found along soft modes.

To test the efficiency of the BGSD method and compare with established methods, the number of force calls required to find processes with barriers less than 0.5 eV in the Al adatom system was determined. For this comparison, we use a scheme for sampling β values between the reactant minimum and a specified maximum, $\beta_{\max}$, with a weight proportional to the probability of finding a new CP or IP. The energy range was split into N_{bins} bins, where each bin i has a probability distribution of being selected, p_i, which is adjusted over the course of the simulation. The probabilities p_i are set initially to $\frac{1}{N_{\text{bins}}}$. For each new saddle search a bin i is selected from the probability distribution and a random value of β is selected from within bin i. The CP or IP found in the search is then determined to be redundant or new, as compared to a history of previous results. The probability of starting a new search in the bin is then adjusted according to

$$p_i \leftarrow p_i(1 \pm \lambda), \tag{9}$$

where the probability is scaled by $1 + \lambda$ if the result is new and $1 - \lambda$ if redundant. The p_i are then normalized. When all CPs and IPs are found in the β range of interest, the probability distribution becomes flat. For the Al system, $\beta_{\max}$ was set to 0.5 eV above the energy of the reactant state, N_{bins} to 10, and λ to 0.1.

The performance of the BGSD method is shown in Fig. 9 (blue line). This plot shows the error in the total rate of escape from the initial state as a function of the average number of force calls for both the BGSD method, averaged

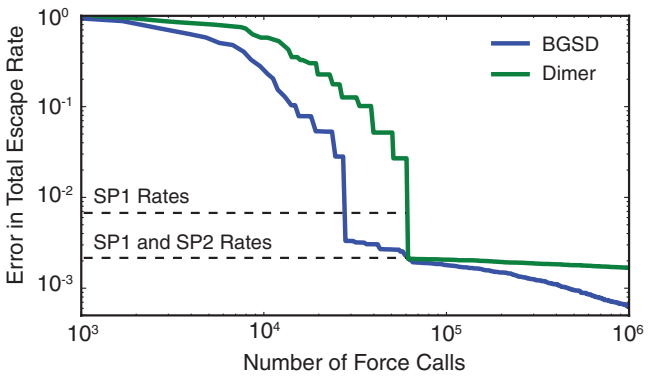

FIG. 9. Comparison of BGSD method versus the dimer method for finding the set of SPs less than 0.5 eV for the Al adatom system.

over ten different runs. The error in the rate is given by

$$k_{\text{Error}} = 1 - \frac{\sum_{i=1}^{n_{\text{SP}}} k_i^{\text{HTST}}}{\sum_{i=1}^{N_{\text{SP}}} k_i^{\text{HTST}}}, \tag{10}$$

where N_{SP} is the number of SPs under 0.5 eV and n_{SP} is the number of SPs currently found by the method.

The efficiency of BGSD was compared to the dimer min-mode following method. Similar local displacements were used to initialize the dimer searches, in which the adatom and all atoms within 3.3 Å were randomly displaced according to a Gaussian distribution with a standard deviation of 0.2 Å. Both the dimer method and the BGSD method used the limited memory Broyden-Fletcher-Goldfarb-Shanno (LBFGS) algorithm for convergence to stationary points.[26] The convergence criterion for the dimer method was $|F| < 0.001$ eV/Å, where F is the force vector. A corresponding final convergence criterion was used for the BGSD method. For $H(r; \alpha, \beta)$ where $\alpha > 0$ convergence was set at $|\nabla H(r; \alpha, \beta)| < 10^{-2}$ eV^2/Å^3. Minimization of $H(r; 0, \beta)$ was terminated when $|H(r; 0, \beta)| < (10^{-3}$ eV/Å$)^2$, which is equivalent to the convergence criterion of the dimer method. The BGSD method can also converge to IPs, for which the optimization was terminated when $|\nabla H(r; 0, \beta)| < 10^{-4}$ eV^2/Å^3 and $H(r; 0, \beta) > (10^{-3}$ eV/Å$)^2$. The comparison in Fig. 9 shows that the BGSD method finds the set of SPs on average with somewhat fewer force calls than the dimer method.

C. Platinum heptamer island

To further test the performance of the BGSD method, a second high-dimensional system was investigated, involving diffusion and rearrangement of a seven-atom Pt-island on a Pt(111) surface. A pairwise Morse potential was used with the same parameters described in Ref. 15. The Pt surface consists of 343 atoms with six layers of 56 atoms, with the seven island atoms on the surface. The bottom three layers of the system were fixed at bulk positions, leaving 175 atoms free to move. The seven lowest energy SPs are shown in Fig. 10. In the two lowest, SP1 and SP2, the entire island shifts to a neighboring site. The next three, SP3–5, involve a two atom slide within the island, and the highest, SP6 and SP7, involve concerted rearrangements of all atoms in the island.

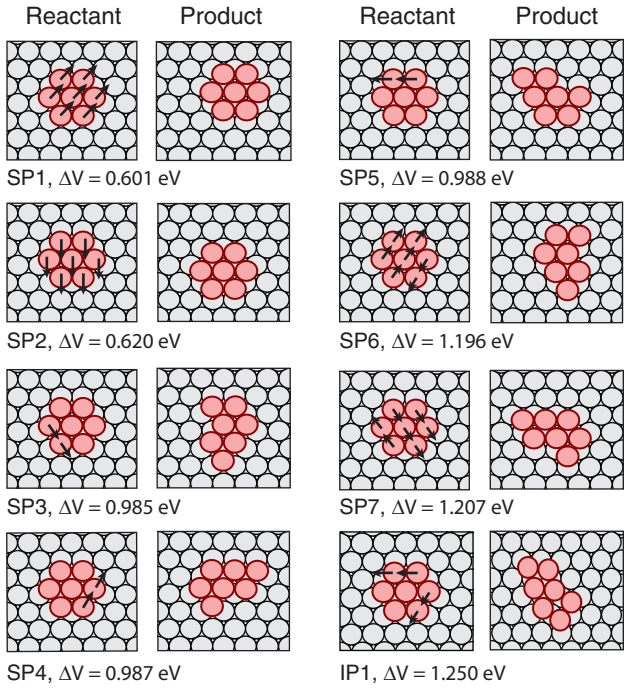

FIG. 10. The seven lowest SPs found from the Pt heptamer island minimum, and the lowest energy IP found.

Figure 11 shows the CP found as a function of β. The BGSD method does a good job finding the low energy SP1–5 as β is increased to 1 eV. Above 1 eV, a large fraction of IPs are found, which is quite different from the Al system. One such IP is shown in Fig. 10. The spurious IPs increase the cost of finding higher energy SPs between 1.0 and 1.5 eV. Additionally, the Pt system is somewhat stiffer than the Al system and required on average twice as many force calls to optimize H and $|\nabla V|^2$.

To compare the efficiency with the dimer method, β values between 0 and 1.5 eV above the minimum were sampled using $N_{\text{bins}} = 15$ and $\lambda = 0.1$. Local displacements were made

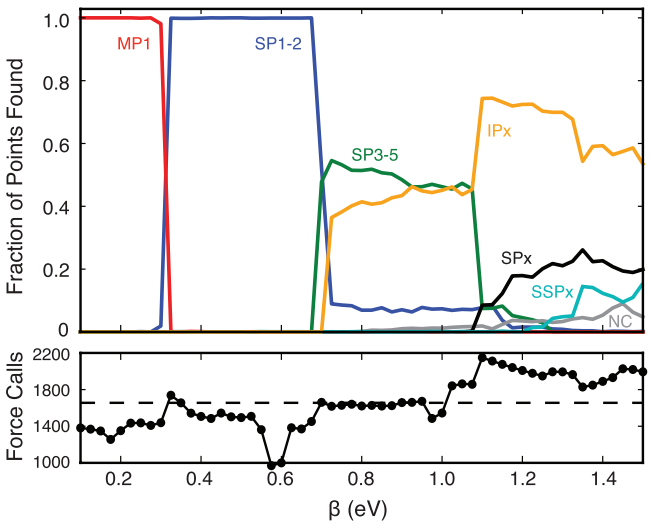

FIG. 11. Fraction of CPs found as a function of β for Pt heptamer diffusion on Pt(111).

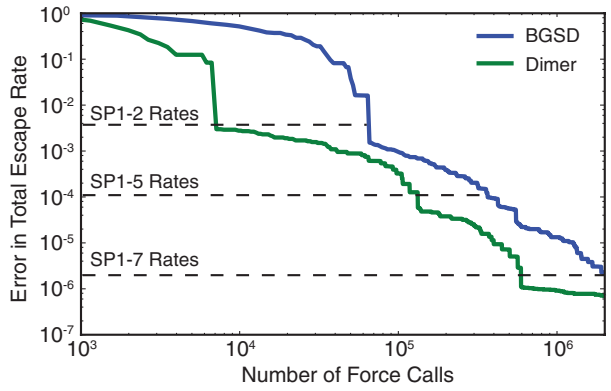

FIG. 12. Comparison of BGSD method versus the dimer method for finding the set of saddle points less than 1.5 eV for the Pt heptamer island system with regards to number of force calls.

within the subspace of a randomly selected atom in the island and all neighbors within 3.3 Å. For the dimer method, Gaussian displacements were made with a standard deviation of 0.1 Å. The same convergence criteria were used as with the Al system. Figure 12 shows that the dimer method is able to find SPs more quickly than BGSD, due to the spurious IPs and slower convergence of each saddle in the Pt system.

IV. CONCLUSIONS

We have introduced a new BGSD method for finding SPs within an energy range of a specified local minimum. This method converts the original potential energy surface, $V(r)$, into a family of objective functions, $H(r; \alpha, \beta)$, which stabilize SPs. As demonstrated by the Al adatom and Pt heptamer island systems, BGSD is competitive in terms of efficiency with the dimer method for finding the set of low energy SPs. Unlike min-mode following methods, the BGSD method finds not only first order SPs but any CPs on the potential energy landscape. Additionally the method can converge to IPs as well as CPs which reduces the efficiency of SP finding. Differences with existing methods, and particularly the ability to target specific energy values, make BGSD a useful addition to the toolkit of SP finding methods.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under Grant No. CHE-1152342 and the Welch Foundation under Grant F-1841. K.P. also acknowledges National Science Foundation support under DMS-1109127. We all would like to thank the Institute of Pure and Applied Mathematics and the University of California at Los Angeles for hosting the long program "Materials for a Sustainable Energy Future" from September 9 to December 13. Computational resources were provided by the Texas Advanced Computing Center.

[1]H. Eyring, J. Chem. Phys. **3**, 107 (1935).
[2]E. Wigner, J. Chem. Phys. **5**, 720 (1937).
[3]C. Wert and C. Zener, Phys. Rev. **76**, 1169 (1949).
[4]G. H. Vineyard, J. Phys. Chem. Solids **3**, 121 (1957).
[5]H. Jónsson, G. Mills, and K. W. Jacobsen, in *Classical and Quantum Dynamics in Condensed Phase Simulations*, edited by B. J. Berne, G. Ciccotti, and D. F. Coker (World Scientific, Singapore, 1998), pp. 385–404.
[6]G. Henkelman and H. Jónsson, J. Chem. Phys. **113**, 9978 (2000).
[7]G. Henkelman, B. P. Uberuaga, and H. Jónsson, J. Chem. Phys. **113**, 9901 (2000).
[8]W. E, W. Ren, and E. Vanden-Eijnden, J. Chem. Phys. **126**, 164103 (2007).
[9]G. Henkelman and H. Jónsson, J. Chem. Phys. **111**, 7010 (1999).
[10]A. Heyden, A. T. Bell, and F. J. Keil, J. Chem. Phys. **123**, 224101 (2005).
[11]J. Kästner and P. Sherwood, J. Chem. Phys. **128**, 014106 (2008).
[12]R. A. Horn and C. R. Johnson, *Matrix Analysis* (Cambridge University Press, Cambridge, 1985).
[13]L. J. Munro and D. J. Wales, Phys. Rev. B **59**, 3969 (1999).
[14]R. Malek and N. Mousseau, Phys. Rev. E **62**, 7723 (2000).
[15]R. A. Olsen, G. J. Kroes, G. Henkelman, A. Arnaldsson, and H. Jónsson, J. Chem. Phys. **121**, 9776 (2004).
[16]G. Henkelman, G. Jóhannesson, and H. Jónsson, in *Progress on Theoretical Chemistry and Physics*, edited by S. Schwartz (Kluwer Academic, New York, 2000), pp. 269–299.
[17]S. Maeda, K. Ohno, and K. Morokuma, Phys. Chem. Chem. Phys. **15**, 3683 (2013).
[18]K. Promislow and B. Wetton, SIAM J. Appl. Math. **70**, 369 (2009).
[19]N. Gavish, J. Jones, Z. Xu, A. Christlieb, and K. Promislow, Polymers **4**, 630 (2012).
[20]J. W. McIver and A. Komornicki, J. Am. Chem. Soc. **94**, 2625 (1972).
[21]G. M. Torrie and G. Valleau, J. Comput. Phys. **23**, 187 (1977).
[22]J. C. Polanyi and W. H. Wong, J. Chem. Phys. **51**, 1439 (1969).
[23]Two Gaussians were added to the LEPS potential. One Gaussian was centered at (2.021,−0.173) with standard deviations 0.1 and 0.35 in the x and y directions, respectively; the second was centered at (0.8, 2.0) with standard deviations of 0.5 and 0.7.
[24]A. F. Voter and S. P. Chen, Mat. Res. Soc. Symp. Proc. **82**, 175 (1986).
[25]J. B. Maronsson, H. Jónsson, and T. Vegge, Phys. Chem. Chem. Phys. **14**, 2884 (2012).
[26]J. Nocedal, Math. Comput. **35**, 773 (1980).